

# An Easy Path to Exascale

**Scott Pakin**

**Applied Computer Science Group**

**Los Alamos National Laboratory**

**25 April 2012**

# Disclaimer

---

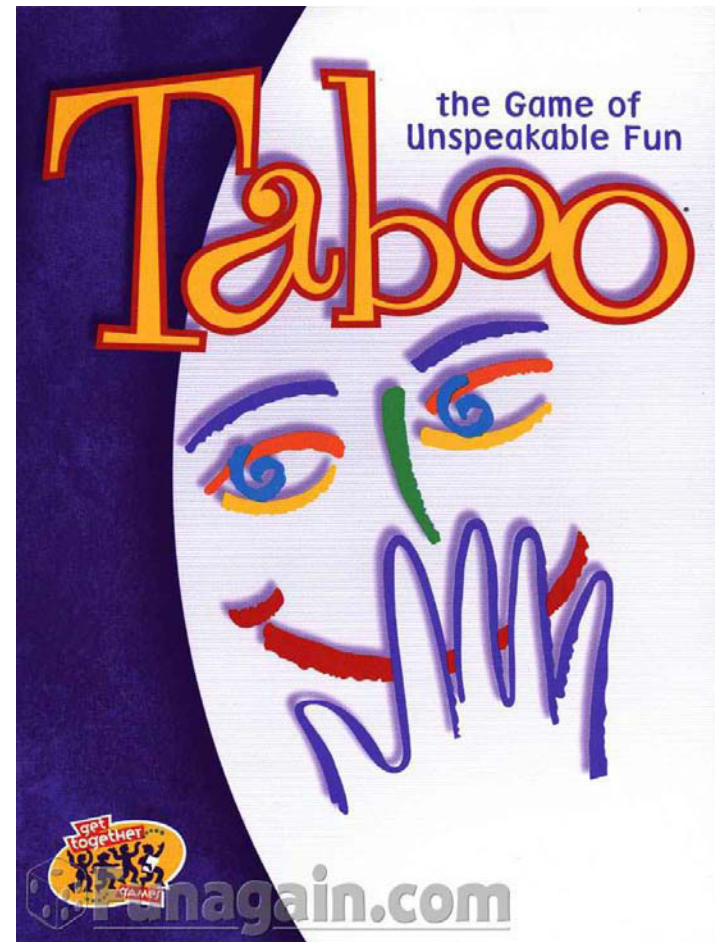
- The opinions expressed herein *definitely* do not reflect the positions of Los Alamos National Laboratory, the National Nuclear Security Administration, or the United States Department of Energy
- Please don't cut off my funding

# Background

- Many of us have been to lots of exascale meetings
- If these meetings forbade the use of the words
  - energy
  - resilience
  - programmability
  - co-design

**the meetings would be a *lot* shorter**

- (but perhaps more fun)



# Exascale Refrains

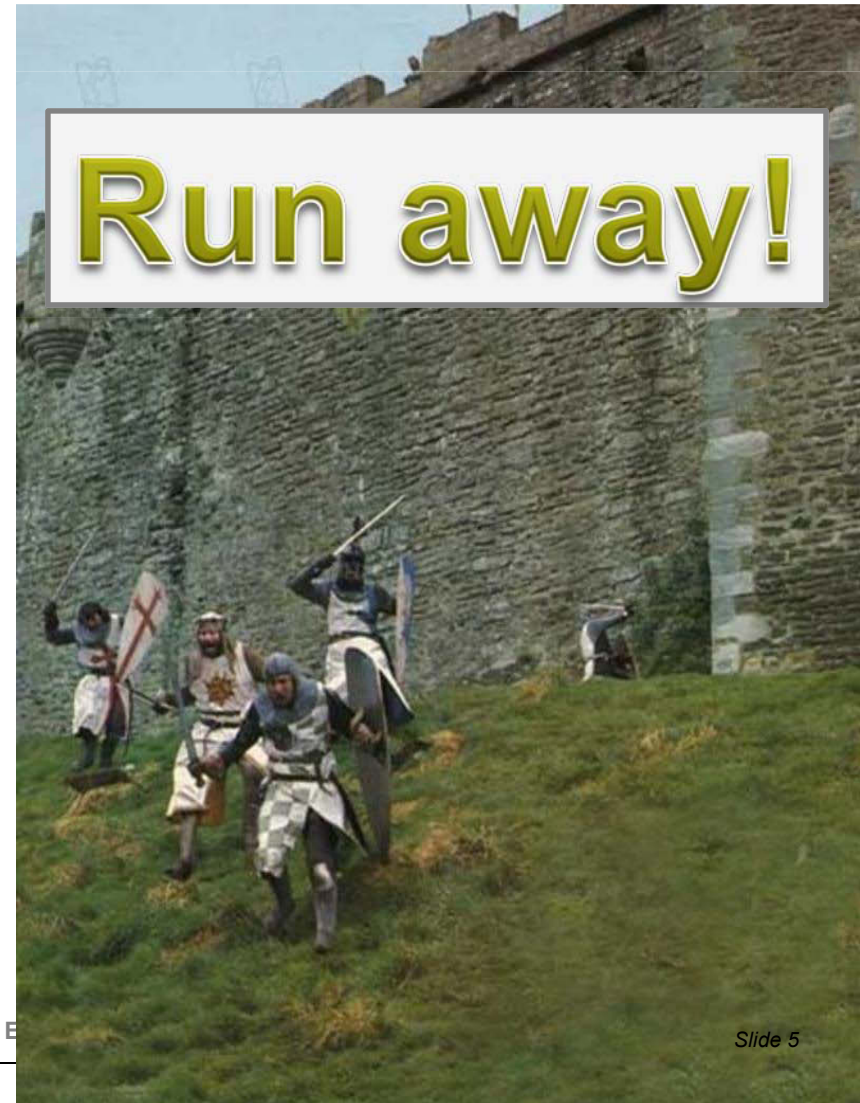
- We don't know how to reduce system power
- We don't have a good way to tolerate frequent faults
- As component heterogeneity increases, programmability decreases

Exascale is hard!



# Solution

- How can we deal with all of the problems of exascale?



# My Proposal

- *What if instead of building an exascale system, we gave each of 1000 researchers his/her own, private, petascale system?*

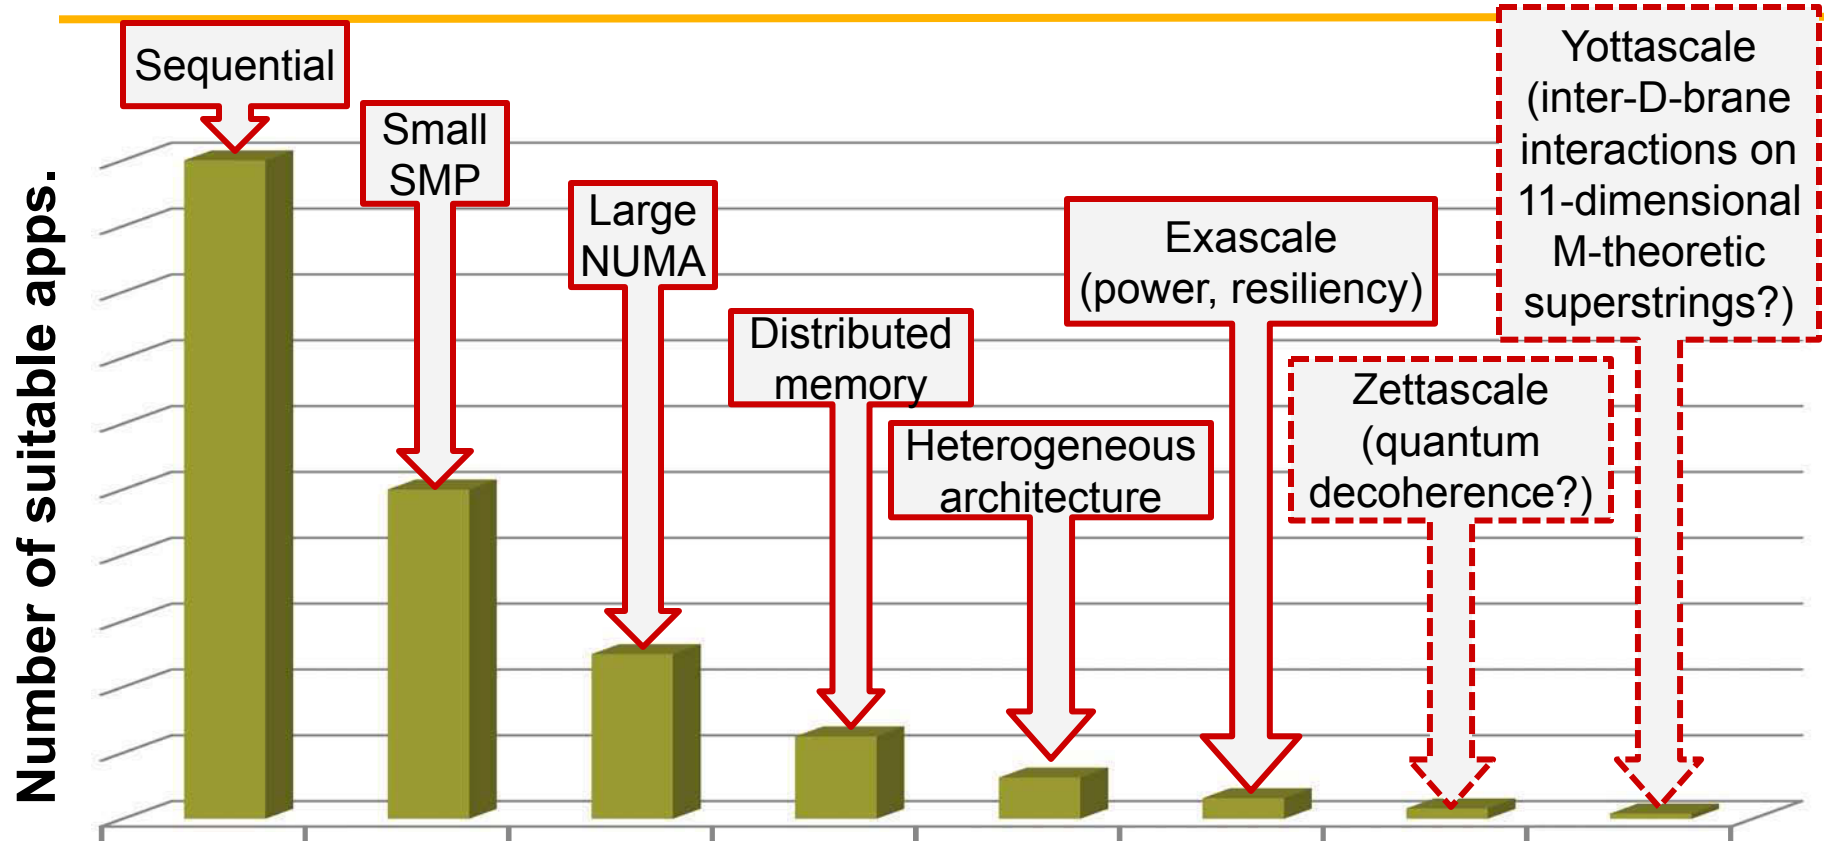


# My Proposal

---

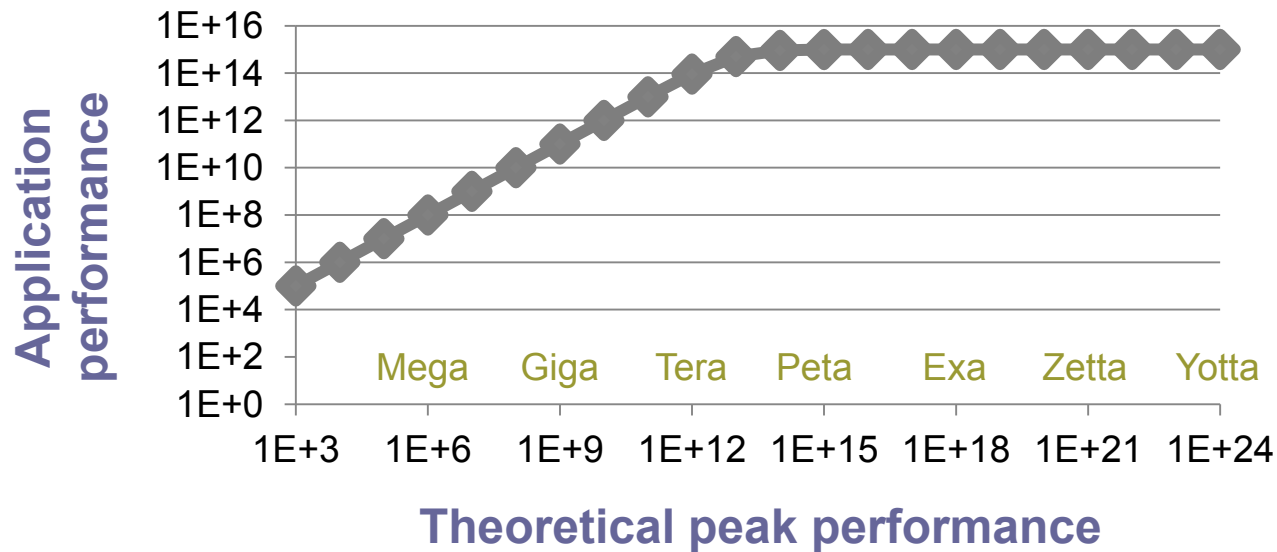
- *What if instead of building an exascale system, we gave each of 1000 researchers his/her own, private, petascale system?*
- Different locations → distribute load across the power grid
  - ✓ Power problem solved
- Fewer components → faults occur less frequently
  - ✓ Resiliency problem solved
- Less complex hardware → can stick with MPI+X (where  $X=\emptyset$ )
  - ✓ Programmability problem solved
- Dedicated systems → no long queues waiting for job to start
  - ✓ Increased productivity
- More aligned with vendor roadmaps → no cajoling needed
  - ✓ Lower cost

# What About Applications that *Require* Exascale?



- Approaching a singularity where we don't *have* any applications (except perhaps LINPACK)

# We May be Facing a Petascale Performance Asymptote



## ■ Parallelism

- By Amdahl's Law, if your application is less than 99.9% parallelized, you'll see at best only petascale performance on an exascale system

## ■ Resilience

- Performing the same work three times gives you 1/3 of the remaining performance

## ■ Power

- To stay within the power budget, some components may have to be turned off or downclocked, lowering performance even further

# Summary

---

- **Exascale computing poses numerous challenges**
  - Energy, resilience, programmability
- **Relatively few applications both require exascale *and* can be programmed to exploit complex exascale hardware**
- **Many applications can take advantage of petascale**
  - Limited by access to machines
  - (Long queue if you want the machine for yourself)
- ***Proposition:* Buy 1000 research teams their own petascale computer**
  - Increases scientists' productivity
  - Why put ourselves through the pain of an exascale before its time?
  - I'll wait until I can borrow an exascale cell phone from my grandkids
- **It's not like there isn't still exciting research to be done at petascale**